

Power Conservation in a Server Cluster

ABSTRACT

5 A system and method for operating a server cluster that includes a set of server devices connected to a local area network (LAN). Each server device maintains a directory of the contents of its file cache. When a decrease in server cluster traffic is detected, a server device on the server cluster is selected for powering down. Prior to powering down a server device, the device's file cache directory is broadcast over the LAN to each of the other server devices on the cluster. If a subsequent request for a file stored in the powered-down server's file cache is received by the cluster, the request is routed to one of the remaining active server devices. This server device then retrieves the requested file from the powered-down server's file cache over the LAN. Prior to broadcasting the file cache directory, pending client requests on the selected server device are completed. The powered-down server may continue to provide power to its 10 NIC and system memory while the processor is deactivated. The server device NIC may include direct memory access capability enabling the NIC to retrieve files from the system memory while the processor is powered-down. 5